

Distributed Applications, Science Grids, and Security

William E. Johnston

Lawrence Berkeley National Lab and NASA Ames Research Center

The vision for “Grids” is to revolutionize the use of computing in science and engineering.

To accomplish this we must make the construction and use of large scale distributed systems that involve diverse resources as easy as using today’s desktop environments.

This ease of use is necessary for R&D agencies like DOE and NASA to routinely address very large simulation and data analysis problems.

The goal of Grids is to provide the *computing infrastructure necessary to routinely build and use dynamically constructed distributed applications.*

Motivating Application Classes

- ***Computational modeling, multi-disciplinary simulation, and scientific data analysis with a world-wide scope of participants*** – e.g. High Energy Physics data analysis, observational cosmology, climate modeling, the National Virtual Air Space
- ***Real-time data analysis and collaboration involving on-line instruments***, especially those that are unique national resources – e.g. LBNL's and ANL's synchrotron light sources, PNNL's gigahertz NMR machines, NASA's turbomachine test cells and Mars sample laboratory, etc.

- ***Generation, management, and use of very large, complex data archives*** that are shared across global science communities – e.g. high energy physics data, Earth environment data (EOS), human genome data
- ***Collaborative, interactive*** analysis and visualization of massive datasets – e.g. DOE's Combustion Corridor project, NASA's air/space frame design data

Grids also have the potential to
provide pools of resources that could be called on in extraordinary / rapid response situations
(such as disaster response) because they can provide:

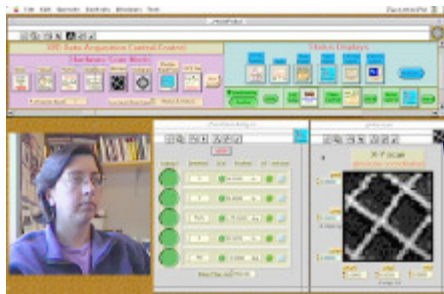
- **common interfaces and access mechanisms**
- **standardized management**
- **uniform user authentication and authorization**

for large collections of distributed resources
(whether or not they normally function in concert).

What are Grids?

Grids are tools, middleware, and services for

- + providing a *uniform look and feel* to a wide variety of computing and data resources
- + supporting *construction, management, and use* of widely distributed application systems
- + facilitating human *collaboration and remote access and operation* of scientific and engineering instrumentation systems
- + *managing and securing* the computing and data infrastructure



Problem Solving Environments and Applications in the Grid Environment



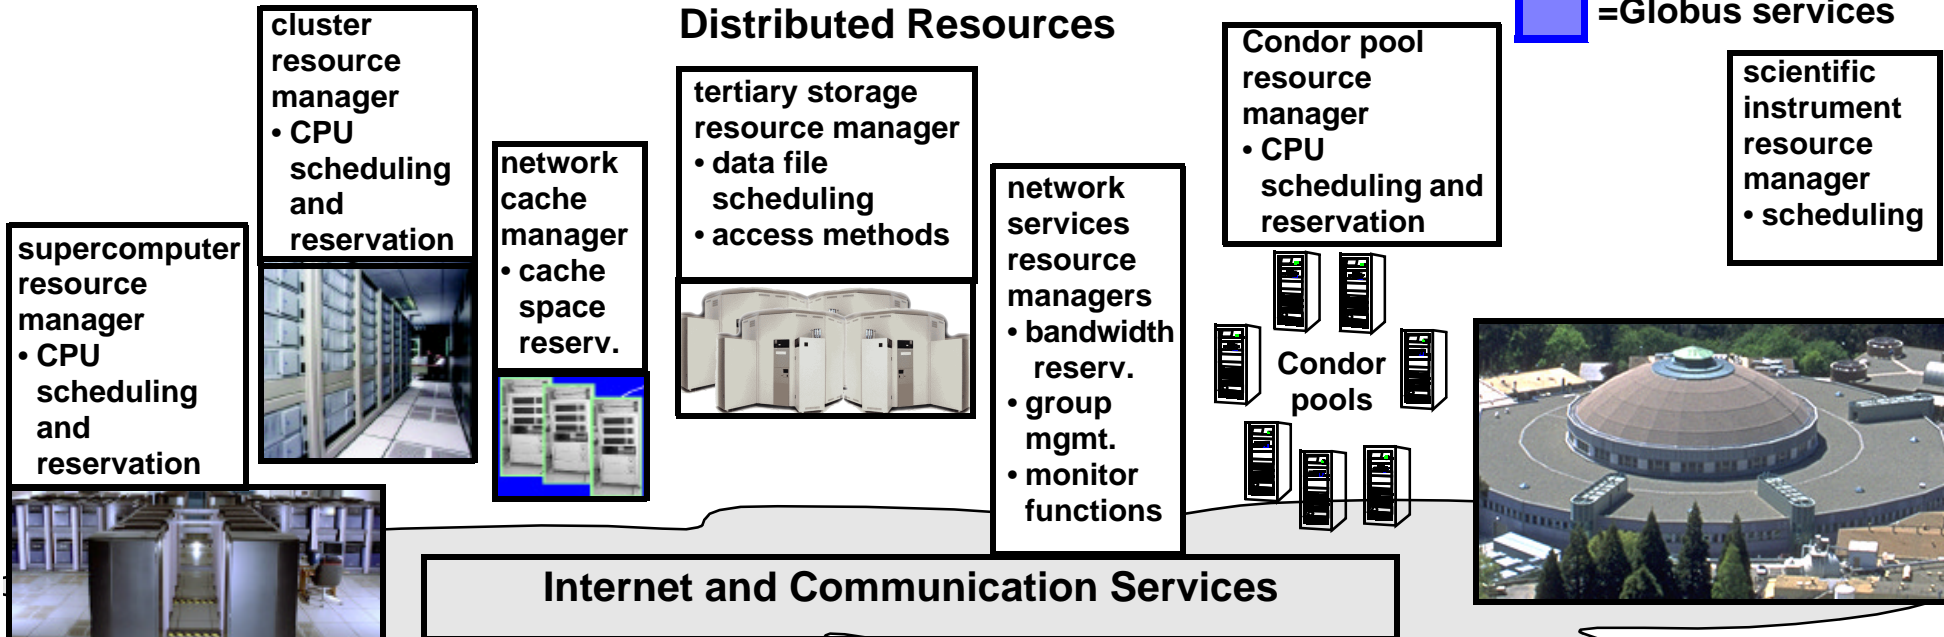
Grid Common Services:

Standardized Services and Uniform Resource Interfaces



=Globus services

Distributed Resources



Internet and Communication Services

The Grid Environment


Software Architecture

of a Grid - upper layers


Problem Solving Environments

- .. Tools to implement the human interfaces
- .. Mechanisms to express, organize, and manage the workflow of a problem solution
- .. Access control
- .. E.g. SciRun [24], Ecce [25], “portals”, WebFlow [26],...

Applications and Supporting Tools



Application Development and Execution Support Services and Systems



Grid Common Services

Distributed Resources

What Grids Will and Will Not Do

Grids provide common resource access technology and operational services deployed across virtual organizations. This *allows the possibility of sharing resources, but does not automatically permit it*:


- + Local authorization models are not changed by the Grid.
- + Common Grid technology will allow common views of resources and uniform access to resources across institutions, thereby permitting very large application systems to be built if policy permits sharing resources across sites and organizations.

Characteristics of Grid-like Systems

- **Kaiser on-line cardio-angiography system**
- **The MAGIC testbed and Terravision - visualization and management of large, distributed earth sciences data sets**
- **The Virtual National Air Space**

WALDO real-time digital library system and DPSS distributed cache for data cataloguing and storage

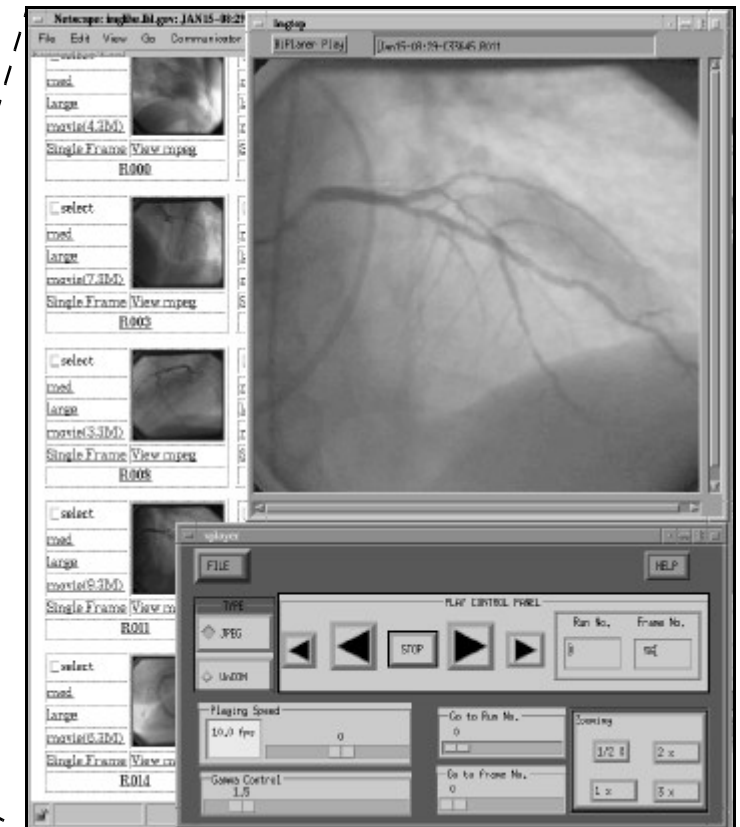
Kaiser San Francisco Hospital Cardiac Catheterization Lab (X-ray video imaging system, ≈ 130 mbit/s, 50% duty cycle 8-10 hr/day)



Tertiary Storage




Compute servers for data analysis and transformation

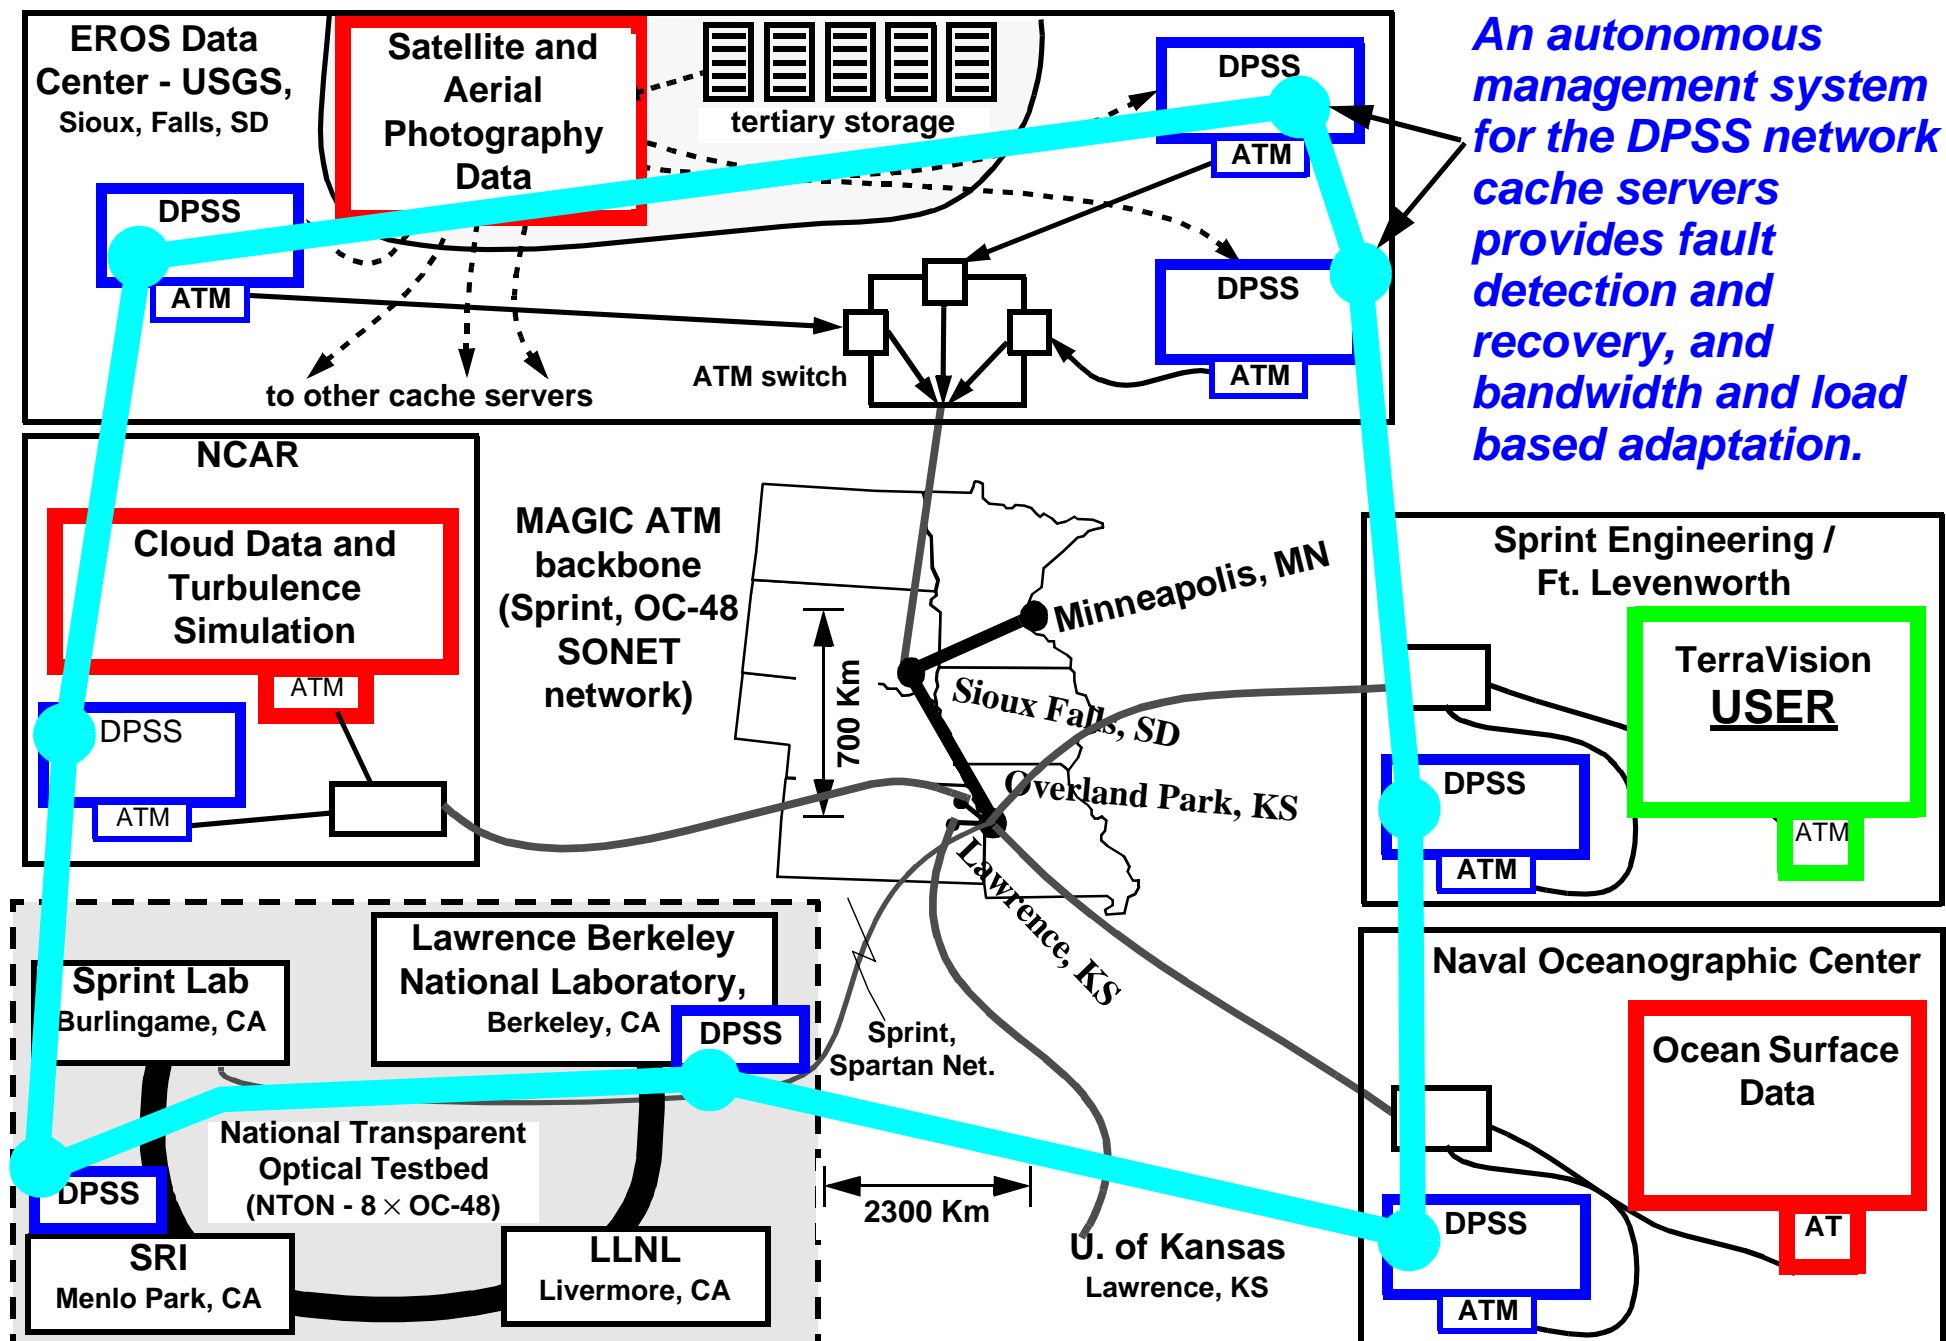


The PSE: Automatically generated user interfaces providing indexed access to the large data objects (the X-ray video) and to various derived data.

Lawrence Berkeley National Laboratory and Kaiser Permanente Health Care
On-line Health Care Imaging Experiment



The Grid Environment



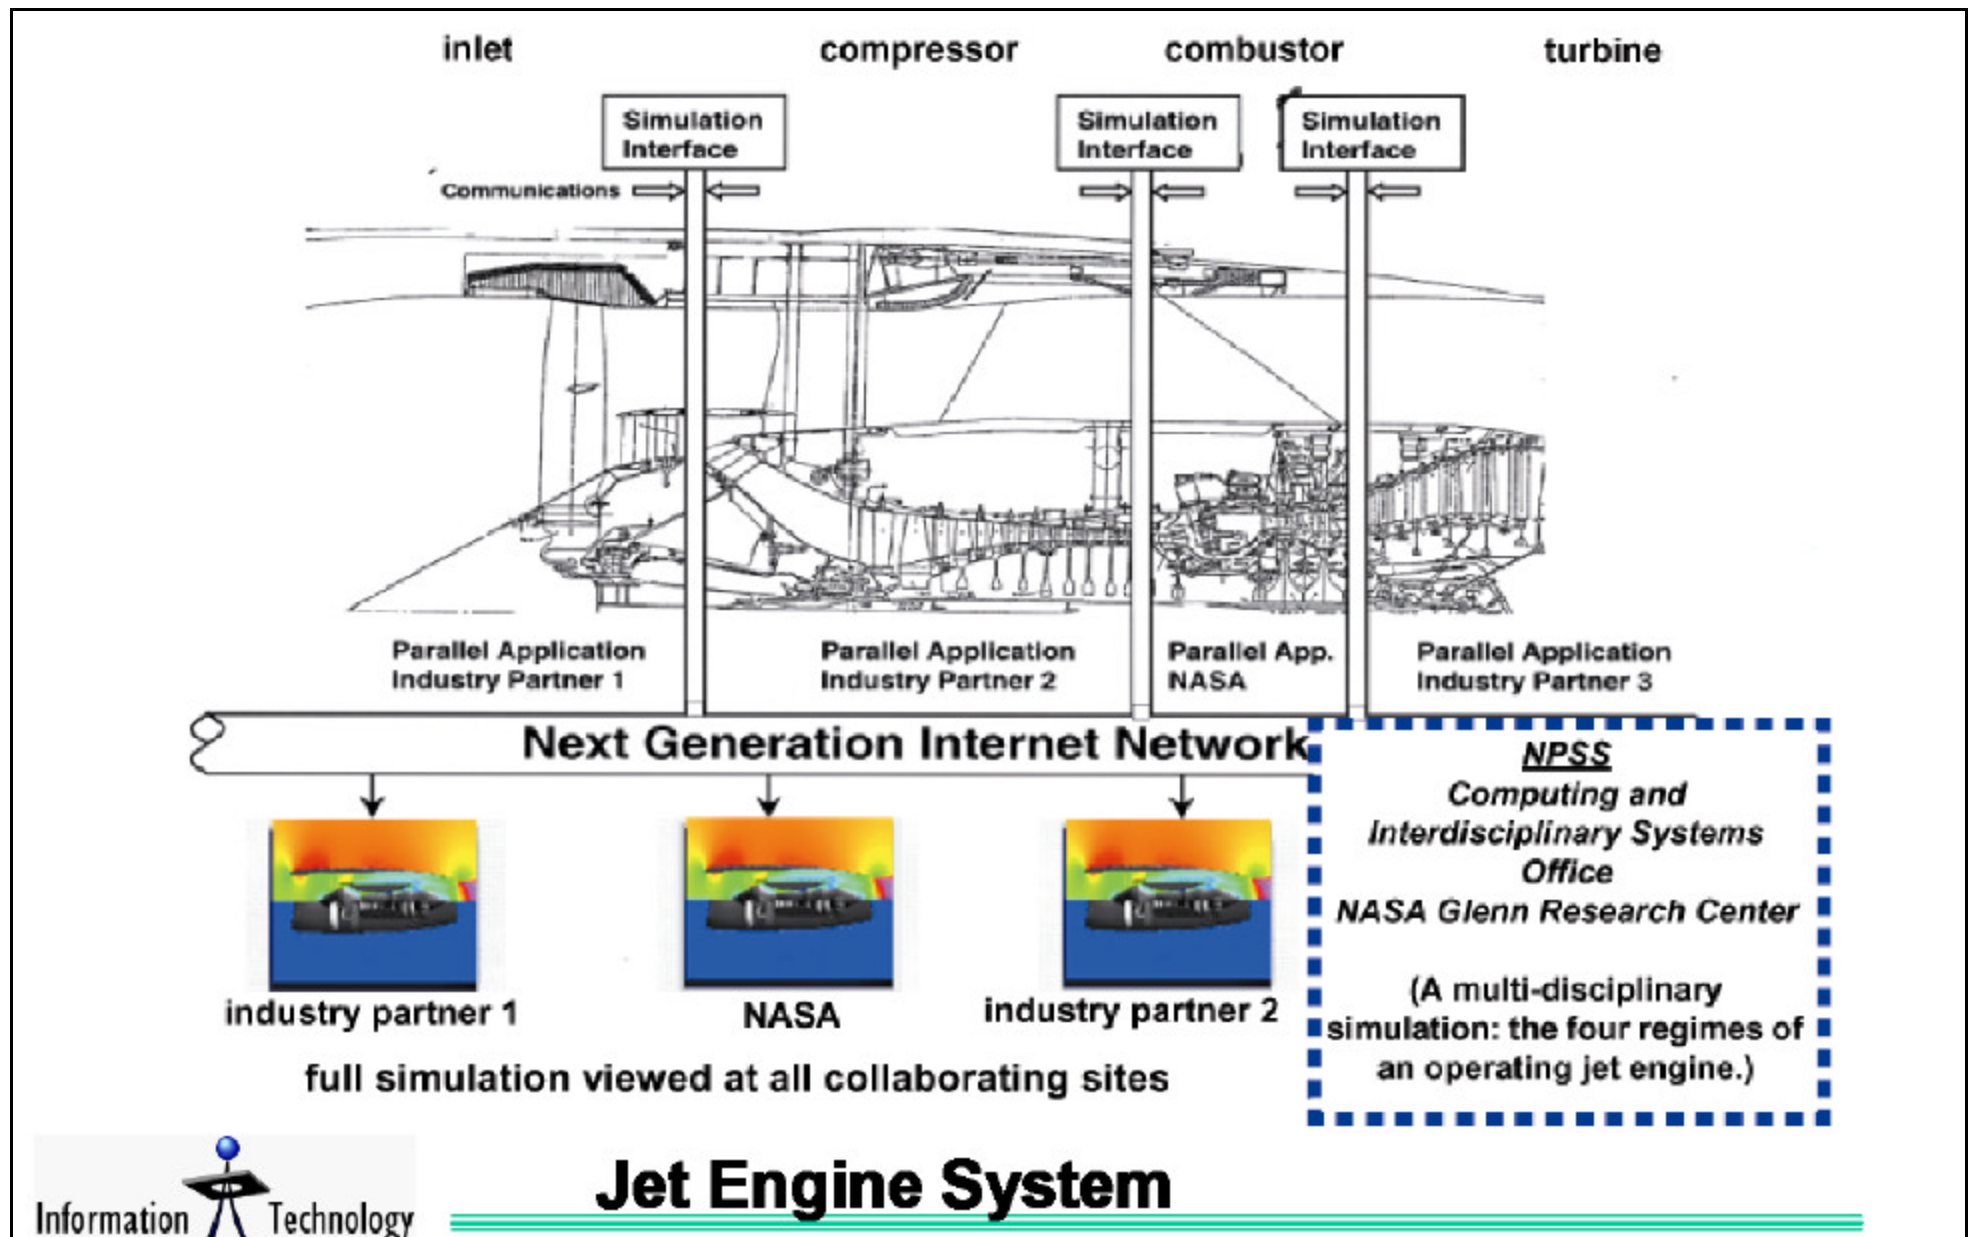
The MAGIC Testbed Distributed Application Environment

A vision for Aviation Safety: Real-time simulation of the entire commercial air space of the country.

(Yuri Gawdiak (VNAS) and Bill McDermott, NASA Ames, John Lytle and Gregory Follen, NASA Glenn (NPSS)).

This vision is being approached through a set of increasingly complex and computationally intensive integrations:


- Component simulations are combined to get a system simulation.



- **Multiple system simulations are coupled to represent pieces of a device.**

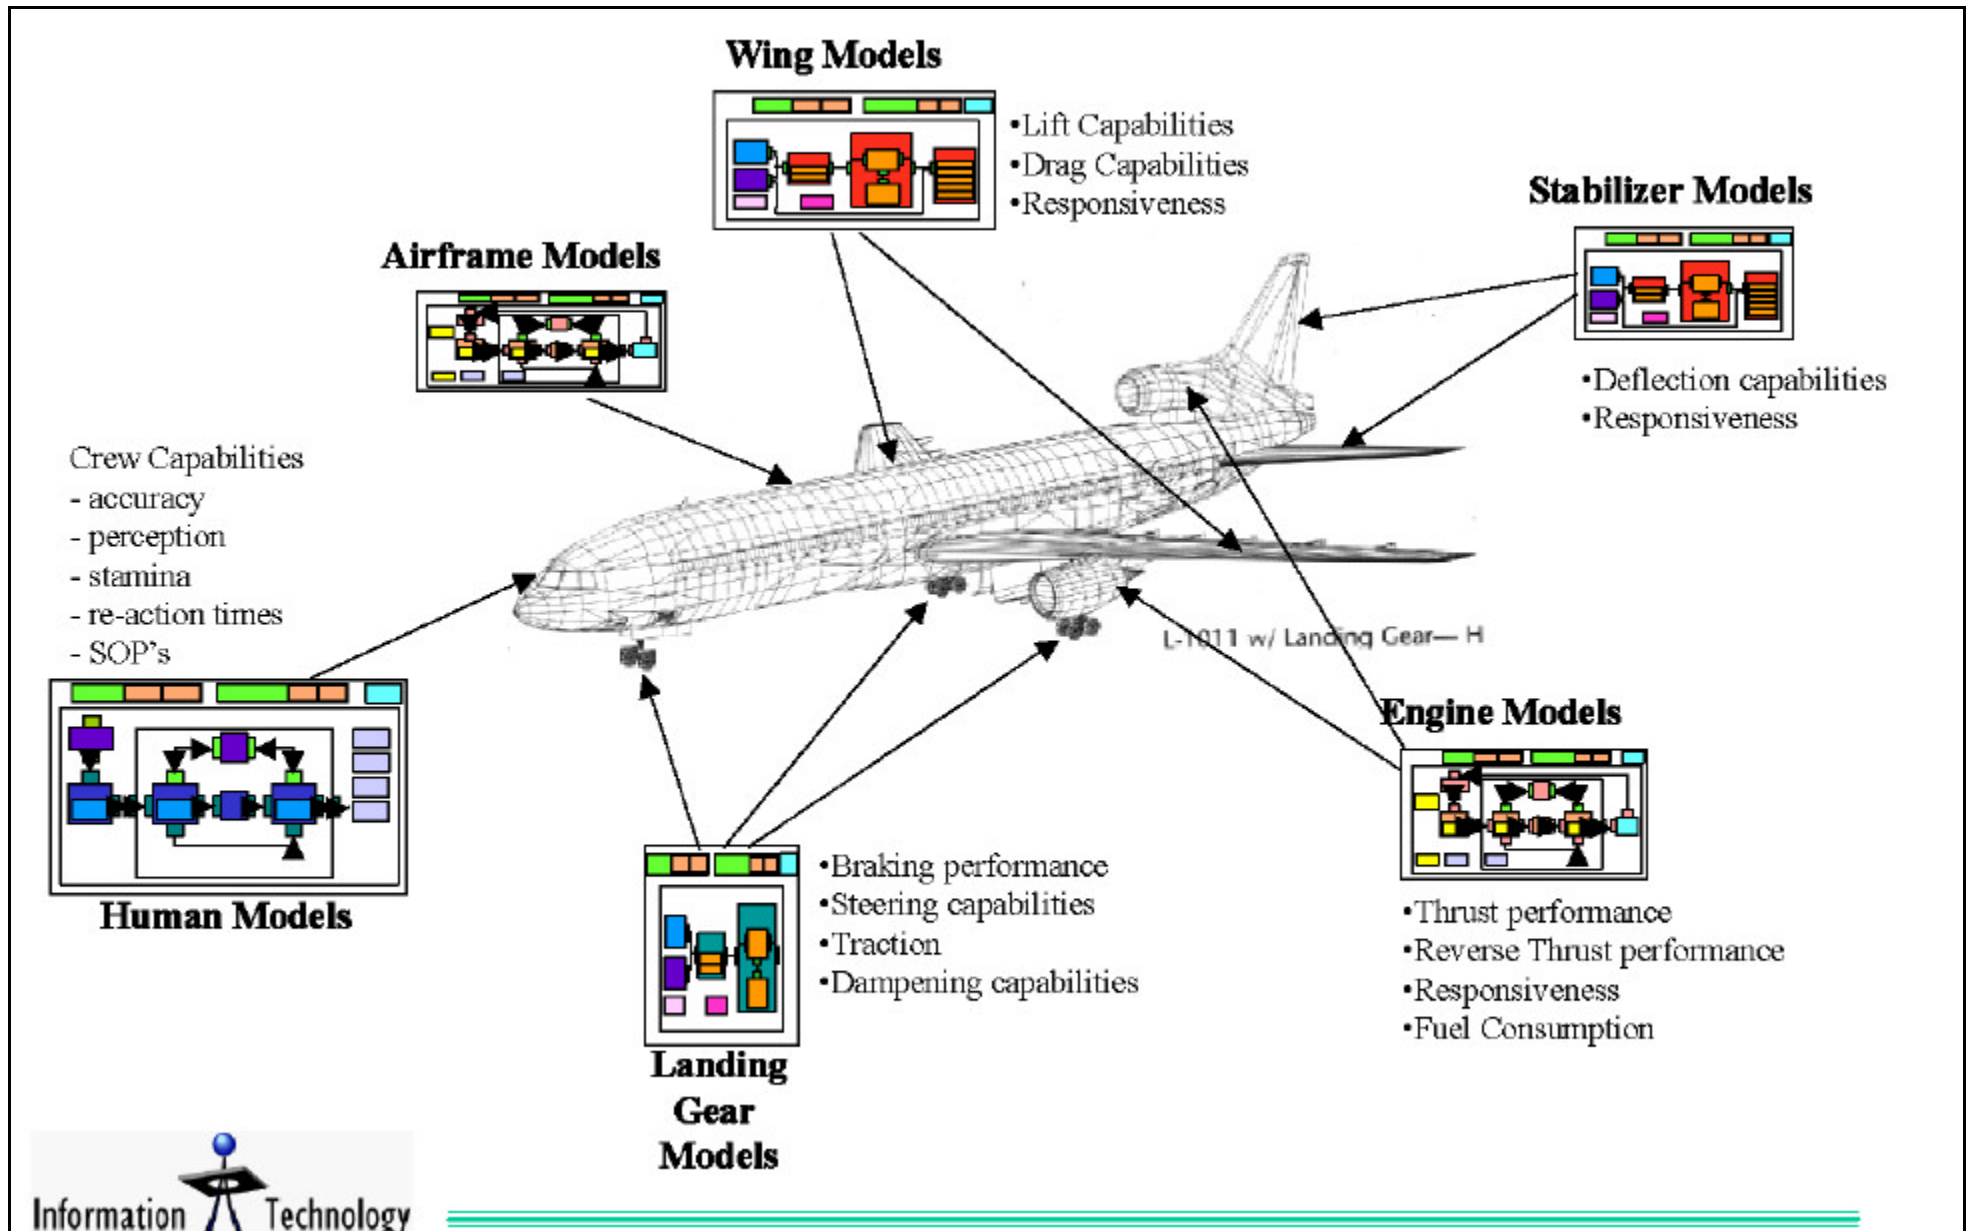


Multi-System Simulation

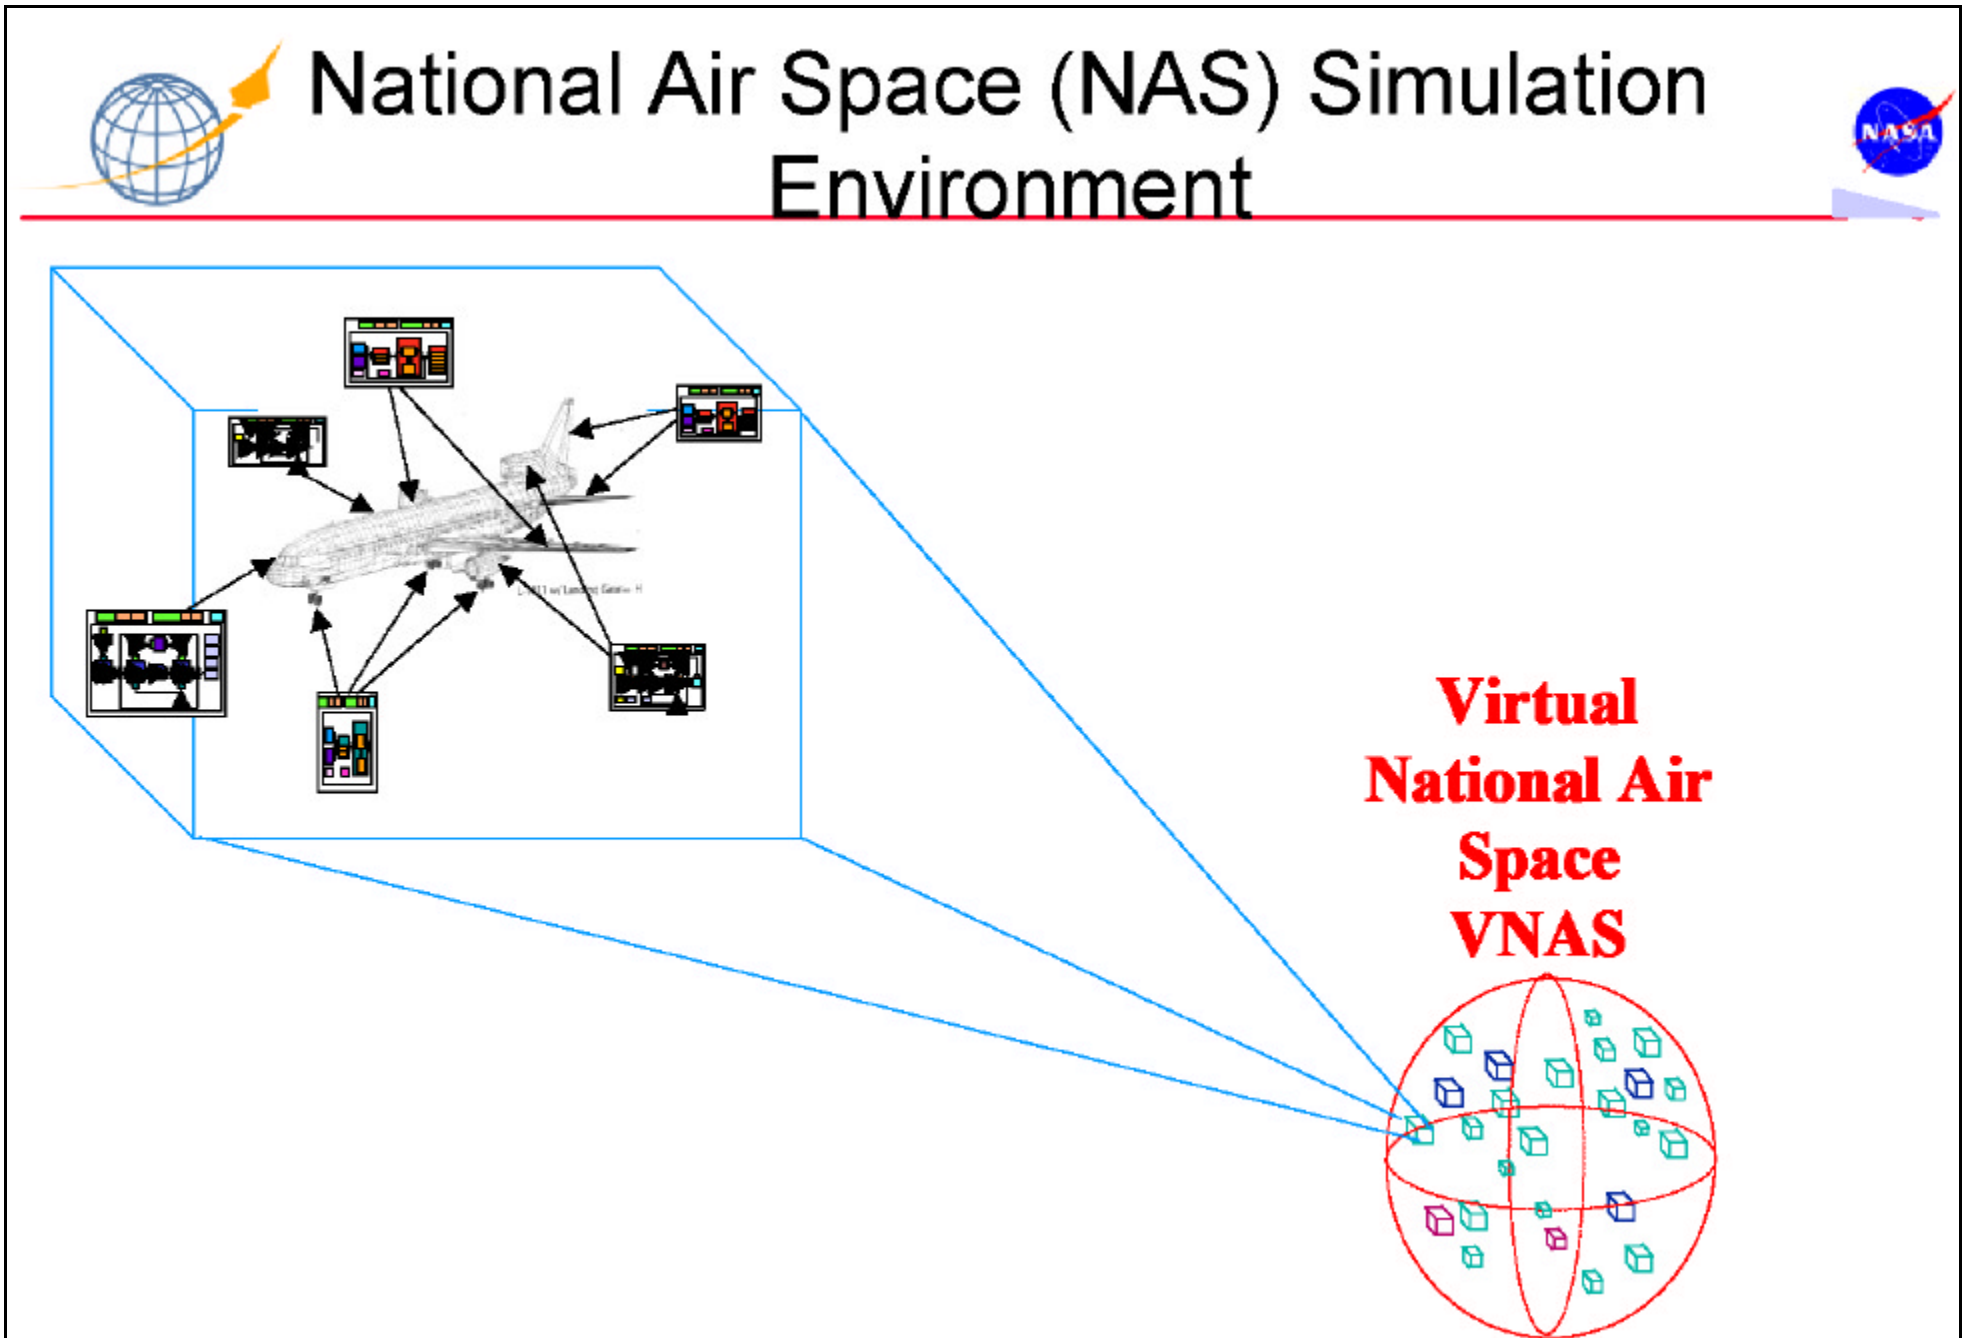


Engine System + Wing System

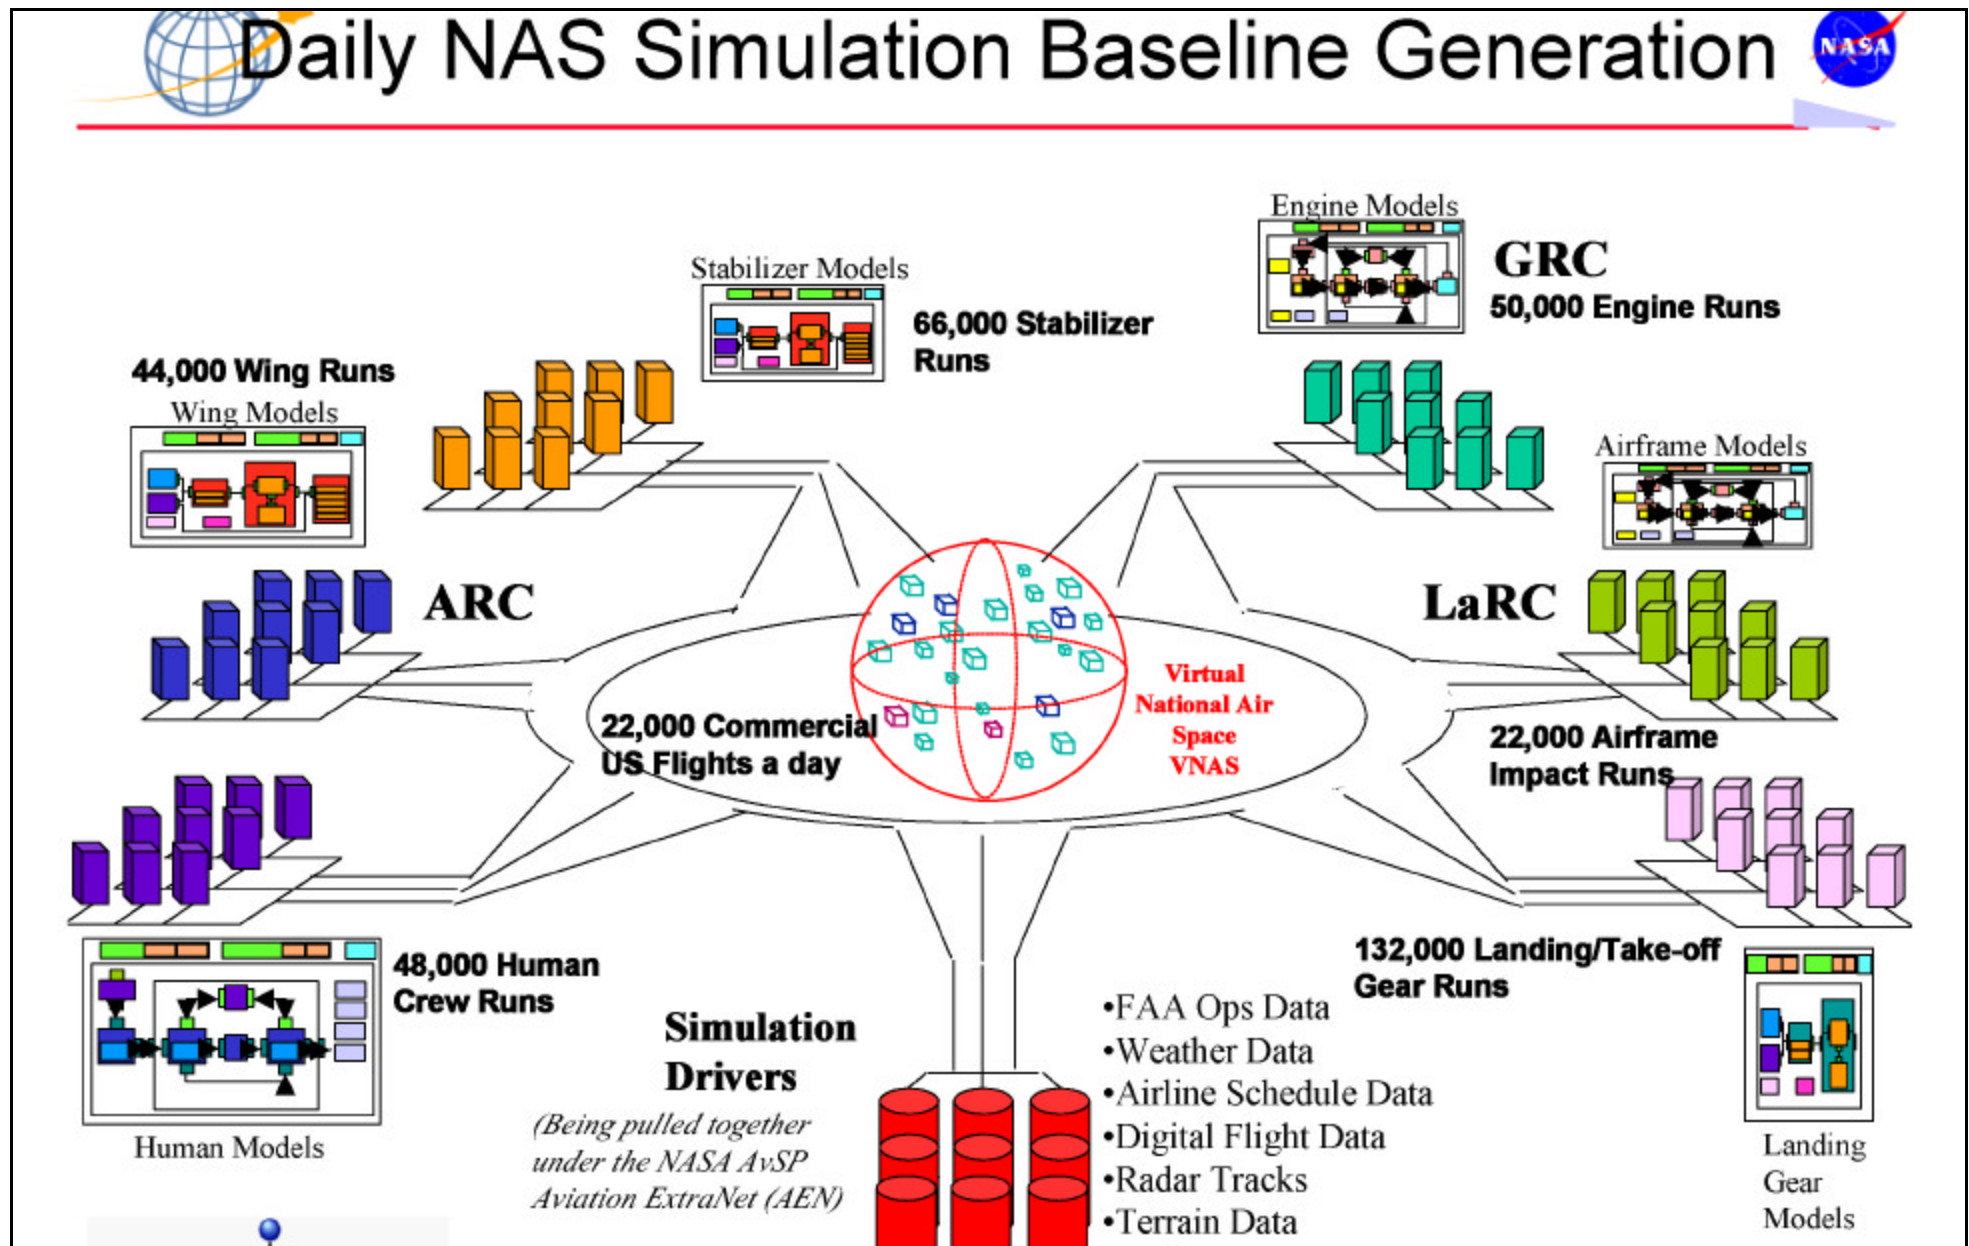
- Whole device simulations are produced by coupling all of the subordinate system simulations.



- Devices are inserted into a realistic environment.




- Devices and environment are combined for operational systems simulation.



Clearly such applications will need to use aggregated computing, data, instrument, and intellectual resources across multiple DOE Labs and NASA Centers, and this will involve, in what are essentially “open” scientific computing environments:

- + locating and coordinating computing and data resources across institutions**
- + multi-institution interactions of many sorts**
- + a shared security model sufficiently well implemented to prevent intrusion, disruption, and theft**



instrument resources

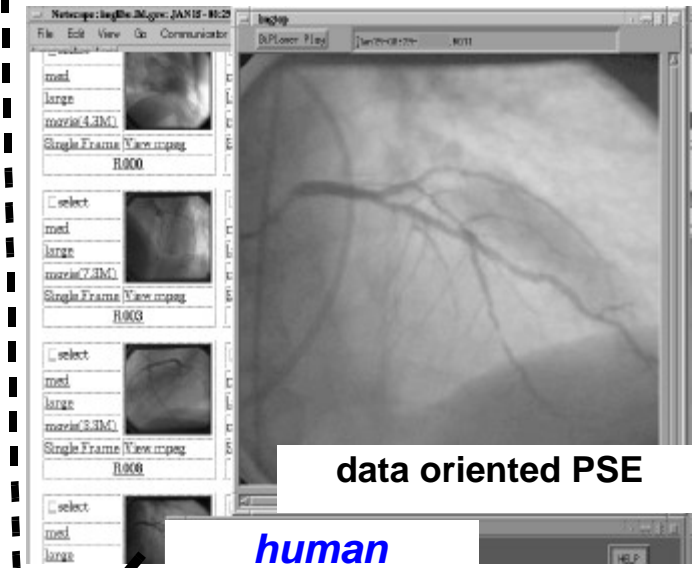


compute resources

resource manager

resource manager

Net3



data oriented PSE

human collaborators

resource manager

resource manager

resource manager

resource manager

Net1

resource manager

storage resources

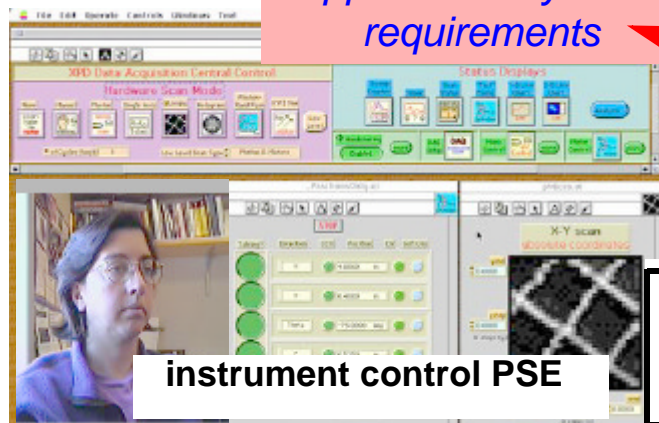
application "system" requirements

resource broker

resource manager

network bandwidth

Net2



instrument control PSE

Real Distributed Applications Require Coordinating Many Resources

Security Aspects of Grids

- **Users are no longer listed in a single central database at a local site, however positive identification to an entity that can provide human accountability will still be required**
 - + **Strong authentication to a globally unique identity**

- **There will be multiple stakeholders for the resources involved in Grid applications who will probably not have a uniform resource use policy: Users will have to be authorized separately for every resource that is incorporated into a Grid application system.**
- + **Strong and flexible policy based authorization and access control**

- **Grid services should not weaken security of local systems, and a security compromise on one platform that is involved in a Grid application system should not propagate via Grid services to other platforms in the system.**

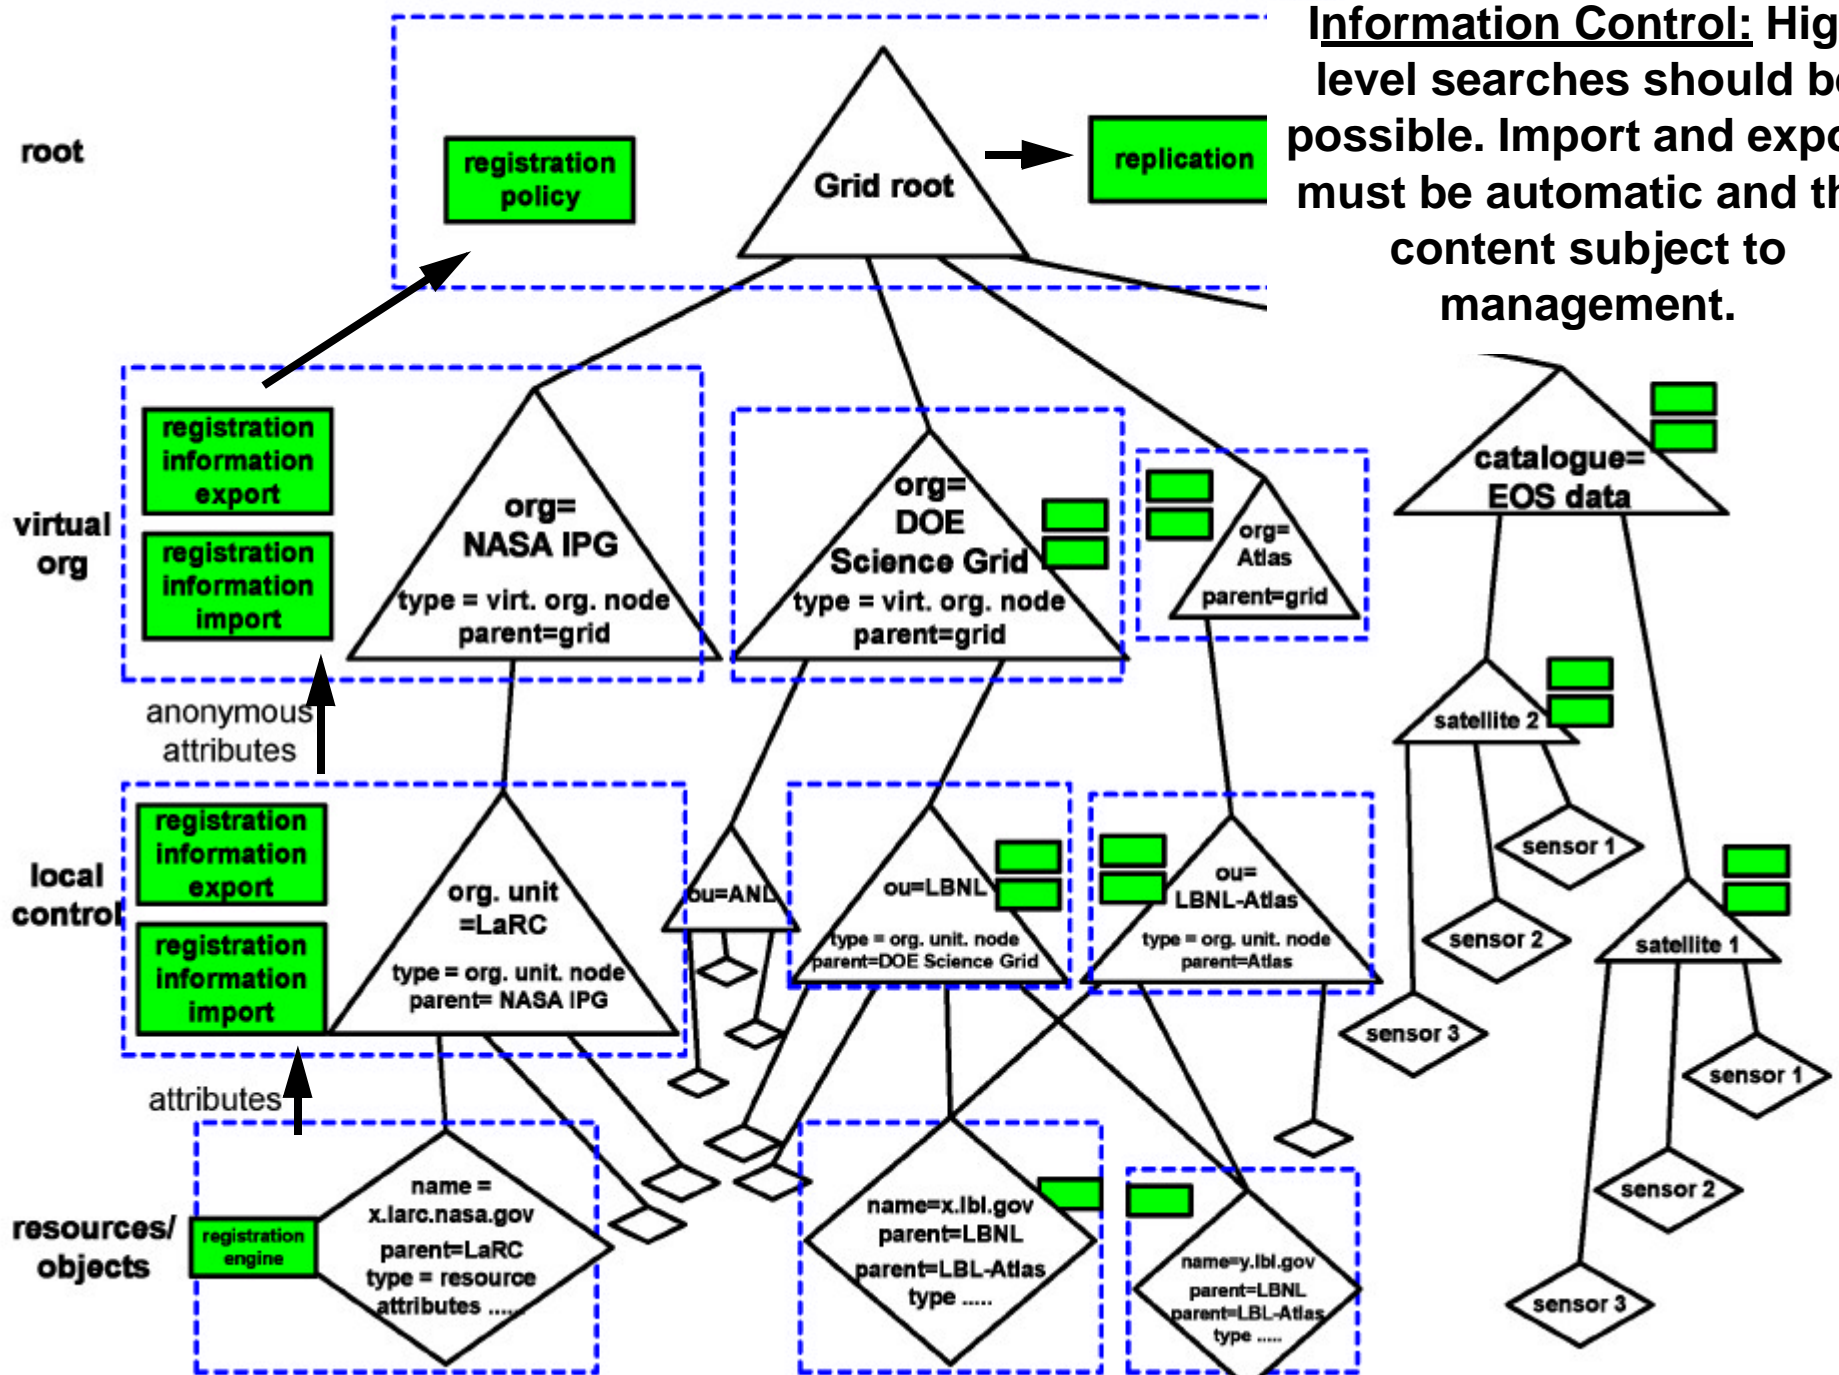
- **Grid users will not have control over the security policy of remote resources (e.g. computing platforms)**
 - + **It may be necessary to “rate” systems on their security, and provide that rating as a system characteristic that may be used in choosing resources from a candidate pool when constructing the resource base for a distributed application.**

- **The Grid Information Service must provide the ability to locate and query information about computing, data, and human resources throughout a potentially global infrastructure.**
- + **The GIS will probably have to have a security model of its own addressing the confidentiality and integrity of the details of site resources while at the same time permitting the sorts information searches needed to assemble application systems from many scattered resources.**

GIS example:

“Within the scope of the Atlas collaboration, return a list of all Sun systems with at least 2 CPUs and 1 gigabyte of memory, and that are running Solaris 2.6 or Solaris 2.7.”

- + Must be able to answer this question, and have control over information propagation.**



Information Control: High level searches should be possible. Import and export must be automatic and the content subject to management.

References and Acronyms

- [1] Globus is a middleware system that provides a suite of services designed to support high performance, distributed applications. Globus provides:
- Resource Management: Components that provide standardized interfaces to various local resource management systems (GRAM) manage allocation of collections of resources (DUROC). All Globus resource management tools are tied together by a uniform resource specification language (RSL).
 - Remote Access: Components that enable remote access to files (GASS and RIO) and executables (GEM).
 - Security: Support for single sign-on, authentication, and authorization within the Globus system (GSI) and (experimentally) authorization (GAA).
 - Fault Detection: Basic support for building fault detection and recovery into Globus applications.
 - Information Infrastructure: Global access to information about the state and configuration of system components of an application (MDS).
 - Grid programming services: Support writing parallel-distributed programs (MPICH-G), monitoring (HBM), etc.

www.globus.org provides full information about the Globus system.

- [2] *The Grid: Blueprint for a New Computing Infrastructure*, edited by Ian Foster and Carl Kesselman. Morgan Kaufmann, Pub. August 1998. ISBN 1-55860-475-8.
http://www.mkp.com/books_catalog/1-55860-475-8.asp

- [3] “Grids as Production Computing Environments: The Engineering Aspects of NASA's Information Power Grid,” William E. Johnston, Dennis Gannon, and Bill Nitzberg. Eighth IEEE International Symposium on High Performance Distributed Computing, Aug. 3-6, 1999, Redondo Beach, California. (Available at <http://www.nas.nasa.gov/~wej/IPG>)
- [4] “Vision and Strategy for a DOE Science Grid” - <http://www.itg.lbl.gov/~wej/Grids>
- [5] See www.nas.nasa.gov/IPG for project information and pointers.
- [6] See <http://www-itg.lbl.gov/NGI/> for project information and pointers.
- [7] The Particle Physics Data Grid has two long-term objectives. Firstly: the delivery of an infrastructure for very widely distributed analysis of particle physics data at multi-petabyte scales by hundreds to thousands of physicists. Secondly: the acceleration of the development of network and middleware infrastructure aimed broadly at data-intensive collaborative science. <http://www.cacr.caltech.edu/ppdg/>
- [8] Tierney, B. Lee, J., Crowley, B., Holding, M., Hylton, J., Drake, F., “A Network-Aware Distributed Storage Cache for Data Intensive Environments”, Proceeding of IEEE High Performance Distributed Computing conference (HPDC-8), August 1999.
- [9] “Real-Time Generation and Cataloguing of Large Data-Objects in Widely Distributed Environments,” W. Johnston, Jin G., C. Larsen, J. Lee, G. Hoo, M. Thompson, and B. Tierney (LBNL) and J. Terdiman (Kaiser Permanente Division of Research). Invited paper, International Journal of Digital Libraries - Special

- Issue on “Digital Libraries in Medicine”. May, 1998. <http://www-itg.lbl.gov/WALDO/>
- [10] MAGIC: “The MAGIC Gigabit Network.” See: <http://www.magic.net>
- [11] TerraVision-2: VRML based data fusion and browsing - www.ai.sri.com/TerraVision
- [12] “A Monitoring Sensor Management System for Grid Environments,” Brian Tierney, Brian Crowley, Dan Gunter, Mason Holding, Jason Lee, Mary Thompson. To appear, HPDC-9, July, 2000. Available at <http://www-didc.lbl.gov/JAMM/>
- [13] A collaborative effort to enable desktop access to remote resources including, supercomputers, network of workstations, smart instruments, data resources, and more - computingportals.org
- [14] “The Data Grid: Towards an Architecture for the Distributed Management and Analysis of Large Scientific Datasets.” A. Chervenak, I. Foster, C. Kesselman, C. Salisbury, S. Tuecke, (to be published in the Journal of Network and Computer Applications).
- [15] “Storage Access Coordination Using CORBA,” A. Sim, H. Nordberg, L.M. Bernardo, A. Shoshani and D. Rotem. Proceedings of the International Symposium on Distributed Objects and Applications. See <http://gizmo.lbl.gov/sm/>
- [16] The Clipper Project: Computational Grids providing middleware that supports applications requiring configurable, distributed, high-performance computing and data resources. See <http://www-itg.lbl.gov/~johnston/Clipper>

- [17] The Grid Forum (www.gridforum.org) is an informal consortium of institutions and individuals working on wide area computing and computational Grids. Current working groups include Security (authentication, authorization), Scheduling and Resource Management, Grid Information Services, Application and Tool Requirements, Advanced Programming Models, Grid User Services and Operations, Account Management, Remote Data Access, Grid Performance
- [18] “New Capabilities in the HENP Grand Challenge Storage Access System and its Application at RHIC” <http://rncus1.lbl.gov/GC/docs/chep292lp1.doc>
“STACS is ... responsible for determining, for each query request, which events and files need to be accessed, to determine the order of files to be cached dynamically so as to maximize their sharing by queries, to request the caching of files from HPSS in tape optimized order, and to determine dynamically which files to keep in the disk cache to maximize file usage.”
- [19] “DeepView: A Collaborative Framework for Distributed Microscopy.” IEEE Conf. on High Performance Computing and Networking, Nov. 1998. See [http://vision.lbl.gov/](http://vision.lbl.gov/projects->collaborative%20computing) (projects -> collaborative computing)
- [20] **Akenti: “Certificate-based Access Control for Widely Distributed Resources,”** Mary Thompson, William Johnston, Srilekha Mudumbai, Gary Hoo, Keith Jackson, Usenix Security Symposium ‘99. Mar. 16, 1999. (See <http://www-itg.lbl.gov/Akenti>)
- [21] GAA: “**Generic Authorization and Access control API**” (GAA API). IETF Draft. http://ghost.isi.edu/info/gss_api.html)

- [22] Storage Resource Broker (SRB) provides uniform access mechanism to diverse and distributed data sources. <http://www.sdsc.edu/MDAS/>
- [23] Condor is a High Throughput Computing environment that can manage very large collections of distributively owned workstations. <http://www.cs.wisc.edu/condor/>
- [24] SCIRun is a scientific programming environment that allows the interactive construction, debugging and steering of large-scale scientific computations. <http://www.cs.utah.edu/~sci/software/>
- [25] Ecce - www.emsl.pnl.gov
- [26] WebFlow - A prototype visual graph based dataflow environment, WebFlow, uses the mesh of Java Web Servers as a control and coordination middleware, WebVM. See <http://iwt.npac.syr.edu/projects/webflow/index.htm>
- [27] “QoS as Middleware: Bandwidth Reservation System Design.” Gary Hoo and William Johnston, Lawrence Berkeley National Laboratory, Ian Foster and Alain Roy, Argonne National Laboratory and University of Chicago. To appear, Eighth IEEE International Symposium on High Performance Distributed Computing, Aug. 3-6, 1999, Redondo Beach, California. (See <http://www-itg.lbl.gov/Clipper/QoS>)